Managing Deadline-constrained Bag-of-Tasks Jobs on Hybrid Clouds

Bo Wang
Department of Computer Science and Technology, Xi'an Jiaotong University
Xi'an, China 710049 &
SKL Computer Architecture, ICT, CAS
wangbo2012@stu.xjtu.edu.cn

Ying Song
SKL Computer Architecture, ICT, CAS
Beijing, China 100190
songying@ict.ac.cn

Yuzhong sun
SKL Computer Architecture, ICT, CAS
Beijing, China 100190
yuzhongsun@ict.ac.cn

Jun Liu
SPKLSTN Lab, Department of Computer Science and Technology, Xi'an Jiaotong University
Xi'an, China 710049
liukeen@mail.xjtu.edu.cn

arXiv:1604.08335v1 [cs.DC] 28 Apr 2016

ABSTRACT

Outsourcing jobs to a public cloud is a cost-effective way to address the problem of satisfying the peak resource demand when the local cloud has insufficient resources. In this paper, we study on managing deadline-constrained bag-of-tasks jobs on hybrid clouds. We present a binary nonlinear programming (BNP) problem to model the hybrid cloud management where the utilization of physical machines (PMs) in the local cloud/cluster is maximized when the local resources are enough to satisfy the deadline constraints of jobs, while when not, the rent cost from the public cloud is minimized. To solve this BNP problem in polynomial time, we proposed a heuristic algorithm. Its main idea is assigning the task closest to its deadline to current core until the core cannot finish any task within its deadline. When there is no available core, the algorithm adds an available PM with most capacity or rents a new VM with highest cost-performance ratio. Extensive experimental results show that our heuristic algorithm saves 16.2%-76% rent cost and improves 47.3%-182.8% resource utilizations satisfying deadline constraints, compared with first fit decreasing algorithm.

Author Keywords

Bag-of-tasks; cloud computing; hybrid cloud; resource management; task scheduling

1. INTRODUCTION

Hybrid cloud, combining a local cloud (private cloud) and a public cloud, is a cost-efficient way to address the problem of insufficient resources in the local cloud when its users have a peak resource demand as the peak load is much larger than average, but transient [1]. Surveyed by the European Network and Information Security Agency (ENISA), most of the small to medium enterprises prefer a mixture of cloud computing models (public cloud, private cloud) [2].

Reducing total capital expenditure on resources is a main objective on hybrid clouds for a provider owning local resources. Generally, using the resources of the local cloud is costless or cheaper, considering that investment costs for the physical infrastructures are "sunk costs", compared with leasing the resources from a public cloud. Thus minimizing the costs for a private cloud provider on a hybrid cloud is the integration of maximizing the resource utilizations of the local cloud and minimizing the rent cost from the public cloud.

There are various researches on scheduling the scientific computing applications on hybrid clouds. A few works focus on minimizing the rent cost with deadline constraints [3–9] or minimizing the makespan [10–12] for scientific computing applications by deciding which tasks should be outsourced to the public cloud. While, these works do not consider how the local cloud/cluster provisions resources, i.e. they do not provide the mapping between physical machines (PM) and provisioned resources in the local cloud.

Only a few hybrid cloud managements [13–15] have coordinated dynamic provisioning and scheduling that is able to cost-effectively complete applications within their respective deadlines. While these works separately scheduled tasks and provisioned resources, i.e., they first decided how many resources, each of which is either a VM in public cloud or a PM in local cloud(s)/cluster(s), used for running tasks and then provisioned the resources from the resource pool, considering that all of the resources are homogeneous.

Different from these existing works, we study on cost-efficiently mapping the tasks to the resources for deadline-constrained Bag-of-Tasks (BoT) jobs, a kind of very common application in the parallel and distributed systems [16, 17], such as parallel image rendering, data analysis, and software testing [18–20], on a hybrid cloud with heterogeneous local

SpringSim-HPC 2016 April 3-6, Pasadena, CA, USA
©2016 Society for Modeling & Simulation International (SCS)

resources. BoT jobs are often composed of hundreds of thousands of independent tasks and are CPU-intensive.

In this paper, we model the task and resource managements of hybrid clouds into a binary nonlinear programming (BNP) model. In this model, the utilization of local resources is maximized while the cost for the resources leased from the public cloud is minimized. As BNP is NP-hard problem [21], we propose a heuristic algorithm to solve this BNP problem. In brief, the contributions of this paper can be summarized as follows:

1. We model the hybrid cloud management which maximizes the utilization of local resources and minimizes the cost for renting the resources of the public cloud for BoT jobs with deadline constraints into a BNP problem.

2. To solve the BNP problem in polynomial time, we propose a heuristic algorithm. The algorithm assigns a task to a core such that the difference between the task's finish time and its deadline is minimum in all assignments between unassigned tasks and cores of used PMs. If there is no such assignment, which means that there is no task can be completed within the deadline by the resources already used, the algorithm adds an available PM with most capacity or leases a VM with most cost-performance ratio from the public cloud when there is no available PM in the local cloud, and assigns tasks to the cores of the added PM/VM as the previous step.

3. We conduct extensive simulation experiments using two real work traces to investigate the effectiveness and efficiency of the proposed algorithm. The experiments results show that our heuristic algorithm saves 16.2%-76% cost for finishing jobs within their respective deadlines and improves 47.3%-182.8% resource utilizations, compared with first fit decreasing (FFD) algorithm.

The rest of the paper is organized as follows. Section 2 presents our model and the heuristic algorithm for solving the model. Section 3 evaluates our work. Section 4 discusses related work and Section 5 concludes this paper.

2. HYBRID CLOUD MANAGEMENT

In a hybrid cloud, as shown in Fig. 1, a task of jobs runs on a core of a PM on the local cloud/cluster or of a VM leased from the public cloud. In this paper, the objective is to cost-efficiently provision resources to tasks and assign the tasks to the resources to meet the complete time within corresponding deadline in the hybrid cloud environment, i.e., to provide the mapping between the PMs or rented VMs and the tasks with minimal cost while fitting deadlines.

2.1. Problem Formulation

We consider a hybrid cloud consisted of a local cloud/cluster and a public cloud. Multiple public clouds in a hybrid cloud can be seen as one big public cloud including the resources provisioned by these public clouds. Table 1 summarizes notations used in this paper.

There are J jobs running on the hybrid cloud. Job i ($i = 1, ..., J$) is composed of T_i independent tasks, $\{t_{i,j} \mid j = 1, ..., T_i\}$. $T = \sum_{i=1}^{J} T_i$ represents the number of all tasks. It needs $r_{i,j}$ resource amounts to complete $t_{i,j}$. Job i must be completed before d_i. Without loss of generality, we assume that $d_1 \leq d_2 \leq \cdots \leq d_J$.

In the local cloud/cluster, there are P PMs. In the public cloud, at most V VMs are rented. We consider a PM or a rented VM as a resource container. Without loss of generality, we assume that the first P containers are PMs and the rest of V containers are VMs. The price per unit time of container k ($k = P + 1, ..., P + V$) is p_k. Container k ($k = 1, ..., P + V$) has N_k cores each of which has r_k capacity. It takes $r_{i,j}/r_k$ time for completing $t_{i,j}$ when running on a core of container k. If all the tasks scheduled to a core can be completed within respective deadlines, respectively, they can run in sequential order by their deadlines in ascending fashion to meet the deadlines. Then the deadline constraints can

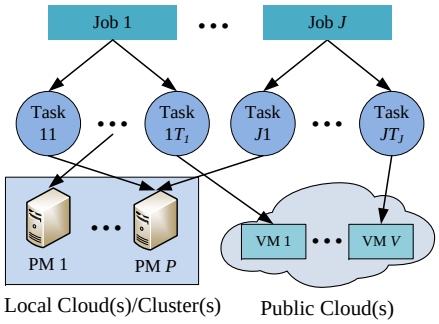

Figure 1. Hybrid cloud environment. A task of jobs is assigned to either a PM in local cloud(s)/cluster(s) or a VM leased from public cloud(s).

Notations	Description
J	The number of jobs running on the hybrid cloud.
T_i	The number of tasks constituting job i.
T	The number of tasks, $T = \sum_{i=1}^{J} T_i$.
$t_{i,j}$	The jth task of job i.
d_i	The deadline of job i.
$r_{i,j}$	The resource amount needed to complete $t_{i,j}$.
P	The number of PMs in the local cloud.
V	The maximal number of VMs leased from the public cloud.
p_k	The price per unit time of container k representing a VM provisioned by the public cloud when $P + 1 \leq k \leq P + V$.
c_k	The cost for renting container k ($P+1 \leq k \leq P+V$).
N_k	The number of cores on container k.
r_k	The capacity of each core on container k.
ε	A small positive constant such that $\varepsilon < \min_{P+1 \leq k \leq P+V} p_k/P$.
$x_{i,j,k,l}$	The binary variable representing whether $t_{i,j}$ is assigned to core l of container k.
N_{RV}	The number of rented VMs.

Table 1. Notations.

be formulated as follow (noticing that $d_1 \leq d_2 \leq \cdots \leq d_J$):

$$\sum_{i'=1}^{i}\sum_{j=1}^{T_{i'}}(x_{i',j,k,l}\cdot\frac{r_{i',j}}{r_k}) \leq d_i,$$
$$\forall i = 1,...,J,\ \forall l = 1,...,N_k,\ \forall k = 1,...,P+V, \quad (1)$$

where the binary variable $x_{i,j,k,l}$ ($i = 1,...,J$, $j = 1,...,T_i$, $k = 1,...,P+V$, $l = 1,...,N_k$) represents whether $t_{i,j}$ is assigned to core l of container k. If so, $x_{i,j,k,l} = 1$, otherwise, $x_{i,j,k,l} = 0$. The left sides of Inequalities (1) represent the finish times of $t_{i,j}$, $i = 1,...,J$, $j = 1,...,T_i$, respectively. The total time using a container is the maximum finish time of the task running on it, $\max_{1\leq l\leq N_k}\sum_{i=1}^{J}\sum_{j=1}^{T_i}(x_{i,j,k,l}\cdot r_{i,j}/r_k)$, $k = 1,...,P+V$. Thus, the costs for leasing VMs from public cloud respectively are

$$c_k = \lceil \max_{1\leq l\leq N_k}\sum_{i=1}^{J}\sum_{j=1}^{T_i}(x_{i,j,k,l}\cdot r_{i,j}/r_k)\rceil \cdot p_k,$$
$$\forall k = P+1,...,P+V, \quad (2)$$

where $\lceil X\rceil$ is the ceiling integer of X, and the utilizations of PMs respectively are

$$u_k = \begin{cases} 0, & \text{if } \sum_{l=1}^{N_k}\sum_{i=1}^{J}\sum_{j=1}^{T_i} x_{i,j,k,l} = 0 \\ \frac{\sum_{l=1}^{N_k}\sum_{i=1}^{J}\sum_{j=1}^{T_i}(x_{i,j,k,l}\cdot r_{i,j})}{N_k\cdot\max_{1\leq l\leq N_k}\sum_{i=1}^{J}\sum_{j=1}^{T_i}(x_{i,j,k,l}\cdot r_{i,j})}, & \text{else} \end{cases},$$
$$\forall k = 1,...,P. \quad (3)$$

The overall utilization of local resources (PMs) is $\sum_{k=1}^{P} u_k/P$.

We formulate the problem of hybrid cloud management as a BNP as follows:

$$\text{Minimize} \quad -\varepsilon\cdot\sum_{k=1}^{P}u_k + \sum_{k=P+1}^{P+V}c_k, \quad (4)$$

subject to:

$$\sum_{k=1}^{P+V}\sum_{l=1}^{N_k}x_{i,j,k,l} = 1,\ \forall j = 1,...,T_i,\ \forall i = 1,...,J, \quad (5)$$

Inequalities (1),

Equations (2),

Equations (3),

$$x_{i,j,k,l} \in \{0,1\},\ \forall j = 1,...,T_i,\ \forall i = 1,...,J,$$
$$\forall l = 1,...,N_k,\ \forall k = 1,...,P+V, \quad (6)$$

The decision variables are $x_{i,j,k,l}$ ($j = 1,...,T_i$, $i = 1,...,J$, $l = 1,...,N_k$, $k = 1,...,P+V$). In the objective function (4), ε is a small constant such that $0 < \varepsilon < p_k/P, \forall k = P+1,...,P+V$, which guarantees that the value of the objective function is negative if there is no VM leased from public cloud while positive if there is at least one rented VM. Therefore, the objective of this model is maximizing the local resource utilization when the resources of the local cloud are enough to complete jobs within respective deadlines while minimizing the cost ($\sum_{k=P+1}^{P+V}c_k$) of rented VMs when the local resources are insufficient. Constraints (5) ensure that each task must be assigned to only one core. Constraints (6) represent the binary requirements for the decision variables. After solving this model, we achieve the task assignments, $x_{i,j,k,l}$ ($i = 1,...,J$, $j = 1,...,T_i$, $k = 1,...,P+V$, $l = 1,...,N_k$), and the renting time for each VM, c_k/p_k ($k = P+1,...,P+V$).

2.2. The Heuristic Algorithm

As BNP is NP-hard [21], we propose a heuristic algorithm to solve the model presented in Section 2.1. in polynomial time. The main idea of the algorithm is to assign the task to a core so that the finish time of the task is closest to its deadline. If there is no enough resource, the algorithm adds an available PM with most capacity or leases a VM with best cost-performance ratio from the public cloud when there is no available PM in the local cloud. The details of the algorithm are described as follows, outlined in Algorithm 1.

If there is no available resource (line 2), the algorithm would add an available PM (lines 3-8) or lease a VM (lines 9-14) from public cloud when there is no available PM in local cloud. The selection principle of an PM is to selecting the PM with most capacity (line 4). The selection principle of rented VM is that the selected VM has the capacity to complete any task when running alone (C1 in line 10) and has best cost-performance ratio (C2 in line 10). If there are multiple types of VMs having same cost-performance ratio, the algorithm selects the VM with minimal price per unit time (C3 in line 10). After selection, the algorithm adds the cores of selected PM (lines 6-8) or rented VM (lines 11-13) to the pool of available cores.

When there are one or more available cores (line 15), the algorithm assigns the unassigned tasks to these cores (lines 15-24). For all of assignments between unassigned tasks and available cores, the algorithm examines the finish times of the tasks and selects the assignment that the difference between the finish time of the task and its deadline is minimal and that the task is finished within its deadline (lines 16-20). If no assignment that the task can be finished within its deadline, there is no available resource for the unassigned task (lines 21-22).

The algorithm repeats the above steps until there is no unassigned task (line 1).

The computing resource consumed by the algorithm are mainly composed of the selection of an available PM or VM and the decision of an assignment between a task and a core. In real world, the numbers of VM instance types and cores in a PM/VM both are a few tens or fewer, thus the selection of a PM or VM and the decision of an assignment are $O(P)$ and $O(T)$, respectively, in time complexity. Therefore, assigning tasks to PMs is $O(T\cdot(P+T))$ in time complexity. We assume that there are N_{RV} VMs leased from public cloud for completing the tasks within their respective deadlines. Assigning tasks to a VM is $O(T)$ in time complexity. Thus, Assigning tasks to the rented VMs is $O(T\cdot N_{RV})$ in time complexity.

Algorithm 1 assigning tasks to PMs or/and rented VMs

$\mathcal{T}$: the set of unassigned tasks, 3-tuples: (an unassigned task, its requirement of resource amount, its deadline);
$\mathcal{C}$: the set of available cores, 4-tuples: (a core, its capacity, the finish time of all tasks assigned to it, the PM/VM containing it);
$\mathcal{PM}$: the set of available PMs, 3-tuples: (an available PM, its core number, the capacity of a core);
$\mathcal{VT}$: the set of available VM instance types, 4-tuples: (an available VM instance type, its core number, the capacity of a core, its price per unit time);
$\mathcal{A}$: the set of assignments, 5-tuples: (a task, the core running it, the PM/VM containing the core, its start time).

Input: $\mathcal{T}; \mathcal{PM}; \mathcal{VT}$
Output: $\mathcal{A}$

while $\mathcal{T} \neq \phi$ **do**
  if $\mathcal{C} = \phi$ **then** */*There is no available resource*/*
    if $\mathcal{PM} \neq \phi$ **then** */*Adding the PM with maximum capacity*/*
      $pm \leftarrow p : (p \in \mathcal{PM}) \wedge (p(2) \cdot p(3) = \max_{p' \in \mathcal{PM}} (p'(2) \cdot p'(3)))$;
      /$p(i)$ is ith element in tuple p.*/*
      $\mathcal{PM} \leftarrow \mathcal{PM} \setminus \{pm\}$;
      for $i = 1$ to $pm(2)$ **do**
        $\mathcal{C} \leftarrow \mathcal{C} \cup \{(\text{new core}, pm(3), 0, pm(1))\}$;
      end for
    else */*Renting an VM with highest cost-performance ratio*/*
      $vm \leftarrow v : (v \in \mathcal{VT}) \wedge \text{C1}(v) \wedge \text{C2}(v) \wedge \text{C3}(v)$;
      /* $\text{C1}(v) : (\forall t)((t \in \mathcal{T}) \wedge (t(2)/v(3) \leq t(3)))$,
      $\text{C2}(v) : v(2) \cdot v(3)/v(4) = \max_{(v' \in \mathcal{VT}) \wedge \text{C1}(v')} (v'(2) \cdot v'(3)/v'(4)))$,
      $\text{C3}(v) : v(4) = \min_{(v' \in \mathcal{VT}) \wedge \text{C1}(v') \wedge \text{C2}(v')} v'(4)$ */
      for $i = 1$ to $vm(2)$ **do**
        $\mathcal{C} \leftarrow \mathcal{C} \cup \{(\text{new core}, vm(3), 0, \text{a new VM instance of type } vm(1))\}$;
      end for
    end if
  else */*Scheduling a task on an available core*/*
    $(task, core) \leftarrow (t, c) : \text{C4}(t, c) \wedge \text{C5}(t, c) \wedge \text{C6}(t, c)$;
    /* $\text{C4}(t, c) : (t \in \mathcal{T}) \wedge (c \in \mathcal{C}) \wedge (t(3) \geq c(3) + t(2)/c(2))$,
    $\text{C5}(t, c) : t(3) - (c(3) + t(2)/c(2)) = \min_{\text{C4}(t', c')} (t'(3) - (c'(3) + t'(2)/c'(2)))$,
    $\text{C6}(t, c) : t(2) = \max_{\text{C4}(t', c') \wedge \text{C5}(t', c')} (t'(2))$ */
    if $\{(task, core)\} \neq \phi$ **then**
      $\mathcal{A} \leftarrow \mathcal{A} \cup \{(task(1), core(1), core(4), core(3), core(3) + task(2)/core(2))\}$;
      $core(3) \leftarrow core(3) + task(2)/core(2)$;
      $\mathcal{T} \leftarrow \mathcal{T} \setminus \{task\}$;
    else */*There is no available resource in $\mathcal{C}$ for any task*/*
      $\mathcal{C} \leftarrow \phi$;
    end if
  end if
end while

Hence, the algorithm is $O(T \cdot (P + T + N_{RV}))$ in time complexity, overall.

3. PERFORMANCE EVALUATION

In this section, we introduce our testbed and experiment design, and then discuss the experimental results.

3.1. Testbed and Experiments Design

We use a 3-month trace collected from the University of Luxemburg *Gaia* cluster system and a 2-month trance collected from NASA Ames iPSC/860, i.e., *UniLu Gaia log* and *NASA iPSC log* in Parallel Workloads Archive [22], to evaluate the performance of our algorithm. We assume that the trace data are the information of tasks running on 1 GHz cores. We set the deadline of each job as α (1, 2, 3, or 4) times of its run time on a 2 GHz core.

Type	CPU Frequency × #cores	#PM
1	2.378 GHz × 8	5
2	2.33 GHz × 8	5
3	2.216 GHz × 4	5

Table 2. The CPU configurations and numbers of PMs used for local resources.

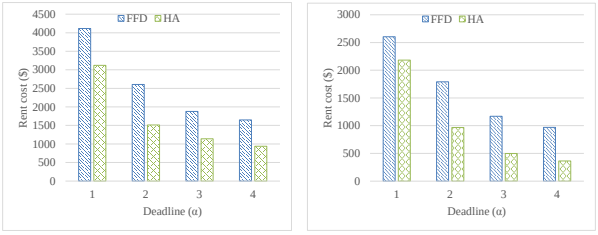

(a) Gaia (b) NASA

Figure 2. The rent costs for finishing Gaia (a) and NASA (b) tasks within their respective deadlines.

The PMs used as the local resources are shown in Table 2. In public cloud, we use a compute optimized instance type, c3.large in EC2 [23], because it has the best cost-performance ratio for CPU-intensive applications, compared with other instances provisioned by EC2. Each VM instance has 2 vCPUs with 2.7GHz. The price of a VM instance is $0.105 per hour (in US east (N. Virginia)).

We use three performance metrics to evaluate task management algorithms:

- ***rent cost***: the cost for leasing VMs from the public cloud;
- ***makespan***: the latest finish time of tasks;
- ***resource utilization***: the overall utilization of used PMs and rented VMs.

3.2. Comparison of Task Management Algorithms

In this section, we compare our Heuristic Algorithm (HA) with First Fit Decreasing (FFD) [24] in both performance and overhead. FFD is assigning the longest task to the first core on which it will fit. The time complexity of FFD is $O(T \cdot (P + N_{RV}))$ [24] if the tasks have been sorted in their sizes. Before running FFD, we sort PMs on their capacities in descending order in advanced.

Performance Comparison

Figures 2, 3, and 4 show the rent costs, the makespans of tasks, and the overall resource utilizations for finishing the tasks, managed by FFD and HA, within their respective deadlines, respectively.

For completing Gaia and NASA tasks within their respective deadlines, respectively, as shown in Fig. 2, HA consumes

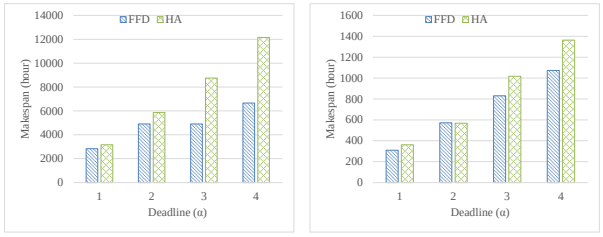

(a) Gaia (b) NASA

Figure 3. The time finishing Gaia (a) and NASA (b) tasks.

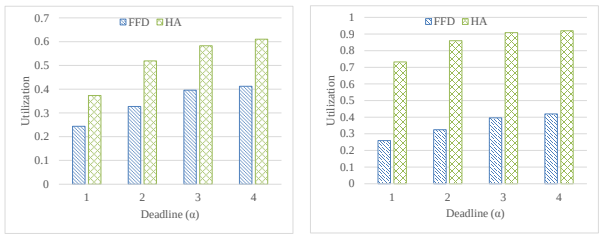

(a) Gaia (b) NASA

Figure 4. The utilization of PMs and VMs running Gaia (a) and NASA (b) tasks.

less 16.2%-76% cost than FFD for leasing VMs from public cloud. While HA postpones the finish time of these tasks to 11.8%-82.6% later, as shown in Fig. 3, compared with FFD, with completing the tasks within their respective deadlines. The less rent cost and longer finish time of HA imply that HA makes better use of resources, compared with FFD. From Fig. 4, we can see that the resource utilization of HA is 47.3%-182.8% more than that of FFD. These can be explained as follows. FFD schedules longest task first, and thus there will be small tasks left to be scheduled after assigning most of tasks. For finishing the left small tasks within their respective deadlines, only a few small tasks are assigned to one rented VM. Therefore, these VMs running only small tasks are under-utilised as rented VMs are charged by unit time, leading to high costs and low utilizations. While HA schedules the task closest to its deadline first, and thus rarely has the problem.

Thus, HA is much better than FFD in minimizing costs and improving resource utilizations in resource and task managements of hybrid clouds.

Figures 2 and 3 show that the rent cost is decreased with increasing the deadline while the finish time increases with deadline. The reason is that tasks are allocated less resources, which leads to longer waiting time of tasks and thus longer time to finish these tasks, if their deadlines are postponed.

Overhead & Scalability

In this section, we experimentally compare HA and FFD in overhead and scalability.

Figure 5 shows the time consumed by HA and FFD running on a Intel(R) Xeon(R) CPU E5410 @ 2.33GHz core. As shown in this figure, HA consumes less time than FFD when the deadline is early ($\alpha = 1/2$), while it consumes

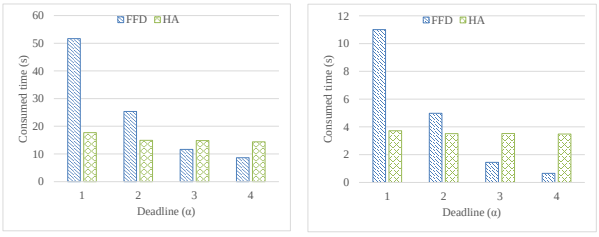

(a) Gaia (b) NASA

Figure 5. The time consumed by task assignments.

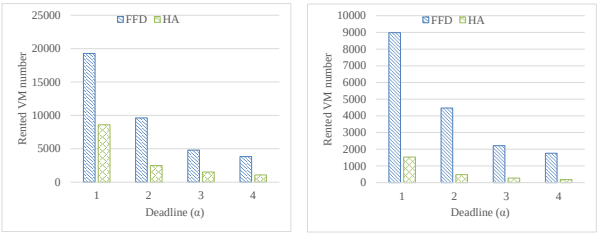

(a) Gaia (b) NASA

Figure 6. The numbers of rented VMs for completing tasks within respective deadlines.

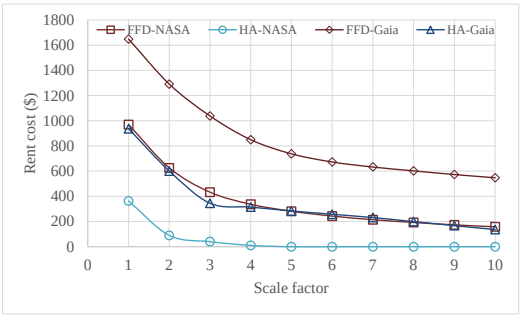

Figure 7. The rent costs decreasing with increasing of PM numbers.

more time than FFD when it is late ($\alpha = 3/4$). The reasons are as follows. The time complexities of HA and FFD are $O(T \cdot (P + T + N_{RV}))$ and $O(T \cdot (P + N_{RV}))$, respectively. When the deadline is earlier, the number of rented VMs is larger. When $\alpha = 1$ or 2, the number of rented VMs used by FFD is larger than 4000, as shown in Fig. 6, which is above 288% more than that by HA and is larger than the number of tasks for both Gaia and NASA traces, and thus the time consumed by FFD is more than HA. When $\alpha \geq 3$, rented VM numbers are smaller than task numbers, therefore, FFD consumes less time than HA.

Next, we examine the scalabilities of HA and FFD on the number of local PMs. We scale the PM resources by a factor ranging from 1 to 10. For example, when the scale factor is 1, 15 PMs described in Table 2 are used for running tasks, while there would be 150 PMs when the scale factor is 10. Figures 7 and 8 show the changes of costs and time consumed by HA and FFD with the scale factor, respectively. We present the results of the case of $\alpha = 4$ here. Other cases have sim-

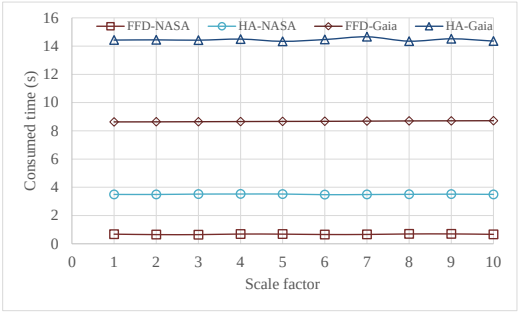

Figure 8. The times consumed by task assignments, stable when the local PM number is increasing.

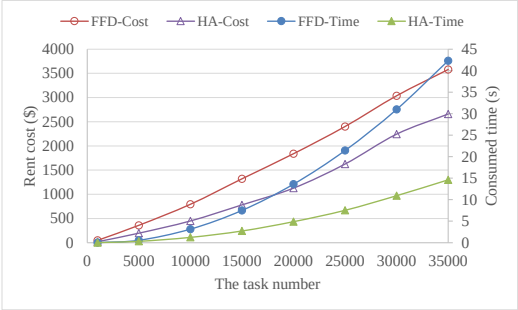

Figure 9. The variations of rented costs and consumed times by task assignments with increasing the task number.

ilar results. As shown in Fig. 7, the cost is decreased with increasing of PM numbers because the resources should be leased from the public cloud are reduced in amount as the amount of local resources increases. From Fig. 7, we can see that HA costs less 24.1%-75.1% and about 20% than FFD for Gaia and NASA traces, respectively, for any scale of local cloud, i.e., HA always consumes less cost than FFD. As shown in Fig. 8, the times consumed by FFD and HA both are stable as the increase of PM number. The reason is that the increase of PM number results in the decreasing of rented VM number, leading to the total number of the PMs and rented VMs ($P + N_{RV}$) having almost no change as the PM number increases.

Now, we examine the scalabilities of HA and FFD on the number of tasks by changing the task scale from 1000 to 35000 in number using Gaia trace. Figure 9 presents the result of the case of $\alpha = 1$. Other cases have similar results. As shown in the figure, the consumed times of HA and FFD both are quadratically increasing with the task number, which is consistent with their respective time complexities. Figure 9 shows that the costs have nearly linear increases with the increase of task number as more resources required for processing more tasks and that HA saves 25.7%-58.1% costs compared with FFD.

These results above indicate that both HA and FFD have good scalability and that HA is better than FFD in minimizing rent cost.

4. RELATED WORK

Outsourcing jobs to a public cloud is a cost-effective way to address the problem of satisfying the peak resource demand when the local cloud/cluster has insufficient resources. There are various researches on scheduling scientific applications on hybrid clouds.

To minimize the cost for leased resources from public clouds, W. Z. Jiang and Z. Q. sheng [25] modelled the mapping of tasks and VMs as a bipartite graph. The two independent vertex sets of the bipartite graph are task and VM collections, respectively. The weight of an edge is the VM cost of a discrete task, i.e. the product of the running time of the task and the cost of the VM per unit time. Then the problem minimizing the cost is to find a subset of the edge set, where the weighted sum of all the edges in the subset is the minimum. The authors used the Hopcroft-Karp algorithm [26] to solve the minimum bipartite match problem. This work does not consider whether a task could be finish within its deadline.

Some existing work studied on minimizing cost with deadline constraints in hybrid clouds. Van den Bossche et al. [3–5] proposed a set of algorithms to cost-efficiently schedule the deadline-constrained BoT applications on both public cloud providers and private infrastructure while taking into account data constraints, data locality and inaccuracies in task runtime estimates. The Hybrid Cloud Optimized Cost (HCOC) scheduling algorithm [6, 7] tried to optimize the monetary execution costs resulting from the public nodes only while maintaining the execution time fitting deadline. HCOC first made an initial schedule using the Path Clustering Heuristic (PCH) algorithm [27] to schedule the tasks to the private cloud and then rescheduled some tasks to the public cloud if the deadline is missed. The algorithm can achieve cost optimization for workflow scheduling. Genez et al. [8] presented an integer linear program model. The numbers of each type of VM instances in both private and public clouds are obtained by solving this model to minimize cost without missing its deadline for a workflow. In this work, the authors considered the VM instances with same type being homogeneous, which is not apply to a private cloud with heterogeneous PMs. For completing the job on time and with minimum cost, Chu and Simmhan [9] first modelled a time-vary spot VM price as a Markov chain and then established a reusable table with three-tuple elements consisted of job compute requirement, deadline constraint of the job and a series of actions with minimal cost based on the price model. The subsequent action was decided by performing a simple table look-up. In this work, they considered that utilizing off local machines does not incur expense.

These above work studied on either resource provisioning or task scheduling. Moreover, most approaches for dynamic provisioning operate in a per-job level, and thus they are inefficient because they fail in consider that other tasks could utilize idle cycles of cloud resources. To address these problems, Aneka [13–15] coordinated dynamic provisioning and scheduling that is able to cost-effectively complete applications within their respective deadlines by considering the whole organization workload at individual tasks level when making decisions and an accounting mechanism to determine

the share of the cost of utilization of public cloud resources to be assigned to each user. While Aneka separately scheduled tasks and provisioned resources, i.e., Aneka first decided how many resources, each of which is either a VM in public cloud or a PM in local cloud/grids, used for running tasks and then provisioned the resources from the resource pool, considering that all of the resources are homogeneous.

Besides minimizing cost, a few work focused on minimizing the makespan of scientific applications by cloud bursting. FermiCloud [10] despatched a VM on the PM that has the highest utilization but still have enough resource for the VM in private cloud. Only when all the resources in private cloud are consumed, VM are deployed on a public cloud. A new VM would be launched in a public cloud only when adding the VM can reduce the average job running time. Kailasam et al. [11, 12] proposed four cloud bursting schedulers whose main ideas are outsourcing a job to a public cloud when the estimated time between now and beginning execution of the job is greater than the estimated time consumed by migrating the job to the public cloud.

These aforementioned work studied on the task and/or resource management with one objective minimizing financial cost for private cloud providers or minimizing the makespan of applications. There are a few work focusing on a balance between two objectives. Taheri et al. [28] proposed a bi-objective optimization model minimizing both the execution time of a batch of jobs and the transfer time required to deliver their required data for hybrid clouds, and used a PSO-based approach to find the relevant pareto frontier. They do not take the finance expenditure into account. V. A. Leena et al. [29] proposed an algorithm for the simultaneous optimization of execution time and cost, in hybrid cloud, by determining whether tasks have to be scheduled to either the private cloud or the public cloud, employing a genetic algorithm. This work does not consider the mapping between VMs and PMs in the private cloud. The VM instances of the same type are homogeneous, which is not true in a heterogeneous data center. Wang et al. [30] proposed a dual-objective multi-dimension multichoice knapsack problem to model the task scheduling with the two objectives of minimizing the cost and minimizing the makespan in hybrid clouds. As the high complexity of solving the problem, the adaptive scheduling algorithm (AsQ) was proposed. AsQ used MAX-MIN strategy [31] to schedule task in private cloud and outsourcing the smallest task to the public cloud when the private cloud has insufficient resources. AsQ allocated the public resource slot with minimal cost to a task, fitting the deadline constraint of the task, considering that using extra public resource slots does not incur extra expense. Hoseiny-Farahabady et al. [32–34] studied on the balance between the makespan and the cost for BoT applications in hybrid clouds. They first established a BNP model with the objective of minimizing the sum of the weighted costs, i.e., the product of cost per unit time of a task and the running time raised to the power of a predefined factor of the task, and then relaxed the model by removing the binary constraints. By Lagrange multiplier method, the relaxed model was solved to get the workload assigned to each resource (PM in the private cloud or VM in public clouds). At last, they used FFD algorithm [24] to assigned tasks to resources so that the total workload of a resource are close to the value obtained from the last step.

Different from all of the above work, we study on cost-efficiently mapping the tasks to the resources for deadline-constrained BoT applications on a hybrid cloud with heterogeneous local resources.

5. CONCLUSION

In this paper, we study on the management of BoT jobs with deadline constraints on hybrid clouds. To minimize the cost of VMs leased from public cloud to complete jobs within their respective deadlines, respectively, we model the task schedule and resource provision into a BNP problem. As the BNP problem is NP-hard, we propose a heuristic algorithm to solve the problem in polynomial time. The main idea of the heuristic algorithm is assigning the task whose finish time is closest to and no later than its deadline if running on the current PM or rented VM. if there is no such assignment, the algorithm add a PM with maximum available capacity or a VM with best cost-performance ratio when there is no available PM. We conduct extensive experiments using two real work traces to investigate the effectiveness and efficiency of the proposed algorithm. The experiments results show that our heuristic algorithm saves 16.2%-76% cost for finishing jobs within their respective deadlines and improves 47.3%-182.8% resource utilizations, compared with first fit decreasing (FFD) algorithm, with good scalability.

ACKNOWLEDGMENT

The authors are grateful to the anonymous reviewers' comments. This work was supported in part by the project of NSFC under grants 61202060, 912183001, 61173112 and 61221062, the National High-Tech Research and Development Program (863) of China under grant 2013AA01A212.

REFERENCES

1. X. Kang, H. Zhang, G. Jiang, H. Chen, X. Meng, and K. Yoshihira, "Measurement, Modeling, and Analysis of Internet Video Sharing Site Workload: A Case Study," in *Web Services, 2008. ICWS '08. IEEE International Conference on*, Sept 2008, pp. 278–285.
2. R. Boutaba and N. L. da Fonseca, "Cloud Architectures, Networks, Services, and Management," *Cloud Services, Networking, and Management*, pp. 1–22, 2015.
3. R. V. den Bossche, K. Vanmechelen, and J. Broeckhove, "Online Cost-efficient Scheduling of Deadline-constrained Workloads on Hybrid Clouds," *Future Generation Computer Systems*, vol. 29, no. 4, pp. 973 – 985, 2013.
4. R. Van den Bossche, K. Vanmechelen, and J. Broeckhove, "Cost-Optimal Scheduling in Hybrid IaaS Clouds for Deadline Constrained Workloads," in *Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on*, July 2010, pp. 228–235.
5. R. V. den Bossche, K. Vanmechelen, and J. Broeckhove, "Cost-Efficient Scheduling Heuristics for Deadline Constrained Workloads on Hybrid Clouds," in *Cloud Computing Technology and Science (CloudCom), 2011 IEEE Third International Conference on*, Nov 2011, pp. 320–327.
6. L. Bittencourt and E. Madeira, "HCOC: a Cost Optimization Algorithm for Workflow Scheduling in Hybrid Clouds," *Journal of Internet Services and Applications*, vol. 2, no. 3, pp. 207–227, 2011.
7. L. Bittencourt, C. Senna, and E. Madeira, "Scheduling Service Workflows for Cost Optimization in Hybrid Clouds," in *Network and Service Management (CNSM), 2010 International Conference on*, Oct 2010, pp. 394–397.

8. T. A. L. Genez, L. F. Bittencourt, and E. R. M. Madeira, "On the Performance-Cost Tradeoff for Workflow Scheduling in Hybrid Clouds," in *Proceedings of the 2013 IEEE/ACM 6th International Conference on Utility and Cloud Computing*, ser. UCC '13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 411–416.
9. H.-Y. Chu and Y. Simmhan, "Cost-Efficient and Resilient Job Life-Cycle Management on Hybrid Clouds," in *Parallel and Distributed Processing Symposium, 2014 IEEE 28th International*, May 2014, pp. 327–336.
10. H. Wu, S. Ren, G. Garzoglio, S. Timm, G. Bernabeu, H. Kimy, K. Chadwick, H. Jang, and S.-Y. Noh, "Automatic Cloud Bursting under FermiCloud," in *Parallel and Distributed Systems (ICPADS), 2013 International Conference on*, Dec 2013, pp. 681–686.
11. S. Kailasam, N. Gnanasambandam, J. Dharanipragada, and N. Sharma, "Optimizing Ordered Throughput Using Autonomic Cloud Bursting Schedulers," *Software Engineering, IEEE Transactions on*, vol. 39, no. 11, pp. 1564–1581, Nov 2013.
12. ——, "Optimizing Service Level Agreements for Autonomic Cloud Bursting Schedulers," in *Parallel Processing Workshops (ICPPW), 2010 39th International Conference on*, Sept 2010, pp. 285–294.
13. R. N. Calheiros, C. Vecchiola, D. Karunamoorthy, and R. Buyya, "The Aneka Platform and QoS-driven Resource Provisioning for Elastic Applications on Hybrid Clouds," *Future Generation Computer Systems*, vol. 28, no. 6, pp. 861 – 870, 2012.
14. C. Vecchiola, R. N. Calheiros, D. Karunamoorthy, and R. Buyya, "Deadline-driven Provisioning of Resources for Scientific Applications in Hybrid Clouds with Aneka," *Future Generation Computer Systems*, vol. 28, no. 1, pp. 58 – 65, 2012.
15. R. Calheiros and R. Buyya, "Cost-Effective Provisioning and Scheduling of Deadline-Constrained Applications in Hybrid Clouds," in *Web Information Systems Engineering - WISE 2012*, ser. Lecture Notes in Computer Science, X. Wang, I. Cruz, A. Delis, and G. Huang, Eds. Springer Berlin Heidelberg, 2012, vol. 7651, pp. 171–184.
16. A. Iosup, M. Jan, O. Sonmez, and D. Epema, "The Characteristics and Performance of Groups of Jobs in Grids," in *Euro-Par 2007 Parallel Processing*, ser. Lecture Notes in Computer Science, A.-M. Kermarrec, L. Boug, and T. Priol, Eds. Springer Berlin Heidelberg, 2007, vol. 4641, pp. 382–393.
17. T. N. Minh, L. Wolters, and D. Epema, "A Realistic Integrated Model of Parallel System Workloads," in *Cluster, Cloud and Grid Computing (CCGrid), 2010 10th IEEE/ACM International Conference on*, May 2010, pp. 464–473.
18. P. of Vision Raytracer Pty Ltd, "Persistence of vision raytracer," http://www.povray.org/, 2013.
19. S. Olivier, J. Huan, J. Liu, J. Prins, J. Dinan, P. Sadayappan, and C.-W. Tseng, "UTS: An Unbalanced Tree Search Benchmark," in *Languages and Compilers for Parallel Computing*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2007, vol. 4382, pp. 235–250.
20. S. Bucur, V. Ureche, C. Zamfir, and G. Candea, "Parallel symbolic execution for automated real-world software testing," in *Proceedings of the Sixth Conference on Computer Systems*, ser. EuroSys '11. New York, NY, USA: ACM, 2011, pp. 183–198. [Online]. Available: http://doi.acm.org/10.1145/1966445.1966463
21. D. Li and X. Sun, *Nonlinear integer programming*. Springer Science & Business Media, 2006, vol. 84.
22. "Parallel Workloads Archive," http://www.cs.huji.ac.il/labs/parallel/workload/, 2015.
23. "Amazon Elastic Compute Cloud (Amazon EC2)," http://aws.amazon.com/ec2/, 2015.
24. J. M. López, J. L. Díaz, and D. F. García, "Utilization Bounds for EDF Scheduling on Real-Time Multiprocessor Systems," *Real-Time Systems*, vol. 28, no. 1, pp. 39–68, 2004.
25. W. Z. Jiang and Z. Q. Sheng, "A New Task Scheduling Algorithm in Hybrid Cloud Environment," in *Cloud and Service Computing (CSC), 2012 International Conference on*, Nov 2012, pp. 45–49.
26. N. Blum, "A New Approach to Maximum Matching in General Graphs," in *Automata, Languages and Programming*, ser. Lecture Notes in Computer Science, M. Paterson, Ed. Springer Berlin Heidelberg, 1990, vol. 443, pp. 586–597.
27. L. F. Bittencourt and E. R. M. Madeira, "A Performance-oriented Adaptive Scheduler for Dependent Tasks on Grids," *Concurrency and Computation: Practice and Experience*, vol. 20, no. 9, pp. 1029–1049, 2008.
28. J. Taheri, A. Y. Zomaya, H. J. Siegel, and Z. Tari, "Pareto Frontier for Job Execution and Data Transfer Time in Hybrid Clouds," *Future Generation Computer Systems*, vol. 37, pp. 321 – 334, 2014.
29. V. Leena and M. Rajasree, "Genetic Algorithm Based Bi-Objective Task Scheduling in Hybrid Cloud Platform," *International Journal of Computer Theory and Engineering*, vol. 8, no. 1, p. 7, 2016.
30. W.-J. Wang, Y.-S. Chang, W.-T. Lo, and Y.-K. Lee, "Adaptive Scheduling for Parallel Tasks with QoS Satisfaction for Hybrid Cloud Environments," *The Journal of Supercomputing*, vol. 66, no. 2, pp. 783–811, 2013.
31. M.-Y. Tsai, P.-F. Chiang, Y.-J. Chang, and W.-J. Wang, "Heuristic scheduling strategies for linear-dependent and independent jobs on heterogeneous grids," in *Grid and Distributed Computing*, ser. Communications in Computer and Information Science, T.-h. Kim, H. Adeli, H.-s. Cho, O. Gervasi, S. Yau, B.-H. Kang, and J. Villalba, Eds. Springer Berlin Heidelberg, 2011, vol. 261, pp. 496–505.
32. M. HoseinyFarahabady, H. Samani, L. Leslie, Y. C. Lee, and A. Zomaya, "Handling Uncertainty: Pareto-Efficient BoT Scheduling on Hybrid Clouds," in *Parallel Processing (ICPP), 2013 42nd International Conference on*, Oct 2013, pp. 419–428.
33. M. HoseinyFarahabady, Y. C. Lee, and A. Zomaya, "Pareto-Optimal Cloud Bursting," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 25, no. 10, pp. 2670–2682, Oct 2014.
34. M. HoseinyFarahabady, Y. Lee, and A. Zomaya, "Randomized Approximation Scheme for Resource Allocation in Hybrid-cloud Environment," *The Journal of Supercomputing*, vol. 69, no. 2, pp. 576–592, 2014.